



DEPLOYMENT GUIDE

DEPLOYING F5 WITH ORACLE'S BEA WEBLOGIC SERVER 10

Table of Contents

Deploying F5 with BEA WebLogic Server 10.0

Prerequisites and configuration notes	1
Configuration example	2

Configuring the BIG-IP LTM for WebLogic Server

Creating the health monitor.....	3
Creating the pool.....	4
Using SSL certificates and keys	6
Creating profiles.....	7
Creating the virtual server.....	14

Configuring the BIG-IP LTM system for WebLogic Portal 17

Creating a health monitor.....	17
Creating the pool.....	17
Using SSL certificates and keys	17
Creating the Profiles.....	17
Creating the Portal virtual server	18

Modifying the WebLogic configuration

Synchronizing the BIG-IP LTM configuration if using a redundant system	20
--	----

Appendix A: Configuring the F5 WebAccelerator module with BEA WebLogic Server

Prerequisites and configuration notes.....	21
Configuration example.....	21
Configuring the WebAccelerator module	22
Connecting to the BIG-IP LTM device.....	22
Creating an HTTP Class profile.....	22
Modifying the Virtual Server to use the Class profile.....	23
Creating an Application	25

Deploying F5 with BEA WebLogic Server 10.0

Welcome to the F5 and Oracle's BEA WebLogic Server deployment guide. F5 provides a highly effective way to optimize and direct traffic for WebLogic Server® 10.0 with the BIG-IP Local Traffic Manager (LTM) and WebAccelerator.

BEA WebLogic Server is at the core of today's most reliable enterprise applications. F5 provides an secure, highly available and scalable application delivery networking device. This strong interoperability and integration provides a solution that delivers unparalleled traffic management functionality for those deploying services and applications on the WebLogic Enterprise Platform.

For more information on BEA WebLogic Server 10, see
www.bea.com/framework.jsp?CNT=index.htm&FP=/content/products/weblogic/server

For more information on the BIG-IP product family, see
www.f5.com/products/big-ip/

Prerequisites and configuration notes

The following are prerequisites for this solution.

- ◆ We recommend the BEA WebLogic server run version 10.0. At a minimum, the WebLogic Server must be running version 5.1.
- ◆ The BIG-IP LTM system must be running version 9.0 or later. We strongly recommend using version 9.4 or later. Some of the examples in this guide use profiles introduced in version 9.4. To use these profiles you must either be running LTM version 9.4 or later, or refer to the Configuration Guide for BIG-IP Local Traffic Management for version 9.4 (available on AskF5), which shows the configuration differences between the base profiles and the optimized profile types.
- ◆ This deployment guide is written with the assumption that you are using the BIG-IP LTM system to offload SSL from the WebLogic Servers. If you are not, simply skip these clearly marked sections.
- ◆ We strongly recommend you read the BEA document: [Load Balancing HTTP Sessions with an External Load Balancer](#)
- ◆ This deployment guide is broken into two sections, one for WebLogic Server, and one for WebLogic Portal.

Product Tested	Version Tested
BIG-IP Local Traffic Manager (LTM)	9.4.4
BEA WebLogic Server	10.0
BEA WebLogic Portal	10.0

Configuration example

Using the configuration in this guide, the BIG-IP LTM system is optimally configured to accelerate and direct traffic to BEA WebLogic Servers. Figure 1 shows a typical configuration with a redundant pair of BIG-IP devices (with an optional WebAccelerator module), a cluster of WebLogic Servers, and a WebLogic administration node. The administration node does not play a part in the F5 configuration.

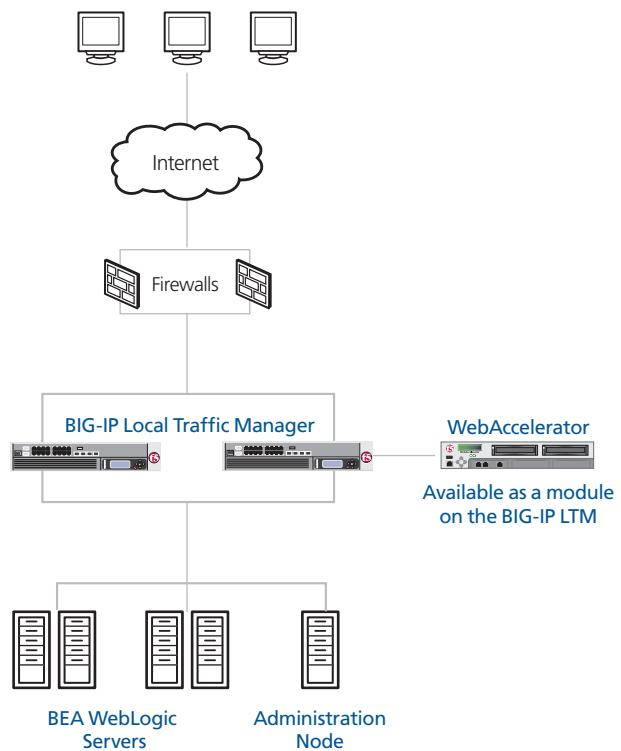


Figure 1 Logical configuration example

This deployment guide is broken up into the following sections:

- *Configuring the BIG-IP LTM for WebLogic Server*
- *Configuring the BIG-IP LTM system for WebLogic Portal*

Configuring the BIG-IP LTM for WebLogic Server

To configure the BIG-IP LTM system for WebLogic Server, you need to complete the following tasks:

- *Creating the health monitor*
- *Creating the pool*
- *Using SSL certificates and keys*
- *Creating profiles*
- *Creating the virtual server*
- *Modifying the WebLogic configuration*

After finishing this section, the following section, *Configuring the BIG-IP LTM system for WebLogic Portal*, on page 17, contains procedures for configuring the BIG-IP LTM for WebLogic Portal Server.

Creating the health monitor

The first step is to set up a health monitor for the WebLogic Servers. This procedure is optional, but very strongly recommended. In our example, we create a simple HTTP health monitor. Although the monitor in the following example is quite simple, you can configure optional settings such as Send and Receive Strings to make the monitor much more specific.

To configure a health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
In our example, we type **wls-http**.
4. From the **Type** list, select **HTTP**.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout. In our example, we use a **Interval of 30** and a **Timeout of 91**.
6. In the **Send String** and **Receive String** sections, you can add a Send String and Receive Rule specific to the device being checked.
7. Click the **Finished** button.
The new monitor is added to the Monitor list (see Figure 2).

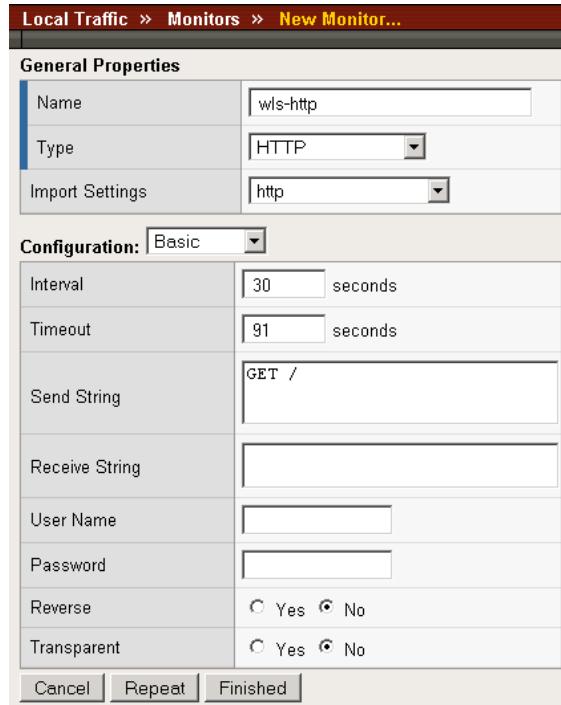


Figure 2 Creating the HTTP Monitor

Creating the pool

The next step is to define a load balancing pool for the WebLogic Servers. A BIG-IP pool is a set of devices grouped together to receive traffic according to a load balancing method. This pool uses the monitor you just created.

To create a pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
The Pool screen opens.
2. In the upper right portion of the screen, click the **Create** button.
The New Pool screen opens.
3. From the **Configuration** list, select **Advanced**.
The Advanced configuration options appear.
4. In the **Name** box, type a name for your pool.
In our example, we use **weblogic-server**.
5. In the **Health Monitors** section, select the name of the monitor you created in the *Creating the health monitor* section, and click the **Add** (<<) button. In our example, we select **wls-http**.
6. In the **Slow Ramp Time** box, type **300**. In our example, we use the Least Connections load balancing method for this pool. We set the Ramp Time in order to ensure that if a pool member becomes

available after maintenance or a new member is added, the Least Connections algorithm does not send all new connections to that member (a newly available member will always have the least number of connections).

If you are not using the Least Connections, Observed, or Predictive load balancing method, skip this step.

7. From the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network).
In our example, we select **Least Connections (member)**.
8. For this pool, we leave the Priority Group Activation **Disabled**.
9. In the New Members section, make sure the **New Address** option button is selected.
10. In the **Address** box, add the first WebLogic server to the pool. In our example, we type **10.133.37.10**
11. In the **Service Port** box, type **7001**.
12. Click the **Add** button to add the member to the list.
13. Repeat steps 10-12 for each server you want to add to the pool.
In our example, we repeat these steps once for the remaining server, **10.133.37.11**.
14. Click the **Finished** button (see Figure 3).

Local Traffic > Pools > **New Pool...**

Configuration: Advanced

Name	weblogic-server
Health Monitors	<div style="display: flex; align-items: center;"> <div style="flex: 1;"> <p>Active</p> <p>wls-http</p> </div> <div style="flex: 1; text-align: right;"> <p>Available</p> <p>oracle10g-portal-https tcp tcp_half_open udp wanjet-eav</p> </div> <div style="flex: 1; text-align: right; margin-top: 10px;"> <> >> </div> </div>
Availability Requirement	All Health Monitor(s)
Allow SNAT	Yes
Allow NAT	Yes
Action On Service Down	None
Slow Ramp Time	300 seconds
IP ToS to Client	Pass Through
IP ToS to Server	Pass Through
Link QoS to Client	Pass Through
Link QoS to Server	Pass Through

Resources

Load Balancing Method	Least Connections (member)
Priority Group Activation	Disabled
New Members	<div style="display: flex; align-items: center;"> <div style="flex: 1;"> <p><input checked="" type="radio"/> New Address <input type="radio"/> Node List</p> <p>Address: 10.133.37.11</p> <p>Service Port: 7001 Select...</p> <p>Add</p> <p>R:1 P:1 10.133.37.10:7001 R:1 P:1 10.133.37.11:7001</p> <p>Edit Delete</p> </div> </div>

Cancel **Repeat** **Finished**

Figure 3 Creating a pool for the WebLogic servers

Using SSL certificates and keys

If you are using the BIG-IP LTM to offload SSL (recommended), you must install a SSL certificate on the virtual server that you wish to use for BEA WebLogic connections on the BIG-IP LTM device. For this Deployment Guide, we assume that you already have obtained an SSL certificate, but it is not yet installed on the BIG-IP LTM system. For information on generating certificates, or using the BIG-IP LTM to generate a request for a new certificate and key from a certificate authority, see the **Managing SSL Traffic** chapter in the *Configuration Guide for Local Traffic Management*.

This section is only necessary if you are using the BIG-IP LTM to offload SSL.

Importing keys and certificates

Once you have obtained a certificate, you can import this certificate into the BIG-IP LTM system using the Configuration utility. By importing a certificate or archive into the Configuration utility, you ease the task of managing that certificate or archive. You can use the Import SSL Certificates and Keys screen only when the certificate you are importing is in Privacy Enhanced Mail (PEM) format.

To import a key or certificate

1. On the Main tab, expand **Local Traffic**.
2. Click **SSL Certificates**. The list of existing certificates displays.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import Type** list, select the type of import (Certificate or Key).
5. In the **Certificate (or Key) Name** box, type a unique name for the certificate or key.
6. In the **Certificate (or Key) Source** box, choose to either upload the file or paste the text.
7. Click **Import**.

If you imported the certificate, repeat this procedure for the key.

Creating profiles

BIG-IP version 9.0 and later use profiles. A *profile* is an object that contains user-configurable settings for controlling the behavior of a particular type of network traffic, such as HTTP connections. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient.

Although it is possible to use the default profiles, we strongly recommend you create new profiles based on the default parent profiles, even if you do not change any of the settings initially. Creating new profiles allows you to easily modify the profile settings specific to this deployment, and ensures you do not accidentally overwrite the default profile.

These profiles use new optimized profiles available in BIG-IP LTM version 9.4 and later. If you are using a BIG-IP LTM version prior to 9.4, the *Configuration Guide for BIG-IP Local Traffic Management* for version 9.4 (available on [AskF5](#)) shows the differences between the base profiles and the optimized profile types. Use this guide to manually configure the optimization settings.

Creating an HTTP profile

The first new profile we create is an HTTP profile. The HTTP profile contains numerous configuration options for how the BIG-IP LTM system handles HTTP traffic. For deployments where WebAccelerator is used, and the majority of users accessing the WebLogic devices are connecting across a WAN, F5 recommends using the **http-acceleration** parent profile (available in versions 9.4.2 and later). This profile uses specific settings to optimize traffic over the WAN. The BIG-IP LTM http-acceleration profile does not have compression enabled by default, because compression is handled by the WebAccelerator. If you are not using the WebAccelerator module, we recommend you use the **http-wan-optimized-compression-caching** parent profile.

If you are not using version 9.4.2, or have other considerations, you can choose the default HTTP parent profile, or one of the other optimized HTTP parent profiles.

To create a new HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **wls-http-opt**.
4. From the **Parent Profile** list, select **http-acceleration**.
5. If you using the BIG-IP LTM to offload SSL, in the **Request Header Insert** row, check the Custom box. In the box, type: **WL-Proxy-SSL: true**.
6. The following is optional, but strongly recommended:
 - a) In the RAM Cache section, click the Custom box for the **URI Caching** row.
 - b) From the URI Caching list, select **URI List**.
 - c) In the **URI** box, type the URI of the login and logout pages, and click the **Exclude** button after each entry. In our example, we type **/groupspace/groupspace.jsp** and **/groupspace/communityFiles/shell/logout.jsp**
7. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.



Figure 4 Configuring the HTTP profile

Creating the TCP profiles

The next profiles we create are the TCP profiles. If most of the BEA WebLogic users are accessing the devices via a Local Area Network, we recommend using the **tcp-lan-optimized** (for server-side TCP connections) parent profile. If the majority of the users are accessing the system from remote or home offices, we recommend using an additional TCP profile, called **tcp-wan-optimized** (for client side TCP connections).

In these profiles, we set the Congestion Control to High Speed, which helps the BIG-IP LTM detect and adapt more quickly to network congestion.

Creating the WAN optimized TCP profile

First we configure the WAN optimized profile. If most of the users are accessing the system over the LAN or other low latency links, you do not need to create this profile.

To create a new TCP WAN optimized profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **wls-wan**.
5. From the **Parent Profile** list, select **tcp-wan-optimized**.
6. In the **Congestion Control** row, click the **Custom** box. From the list, select **High Speed**.

7. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.

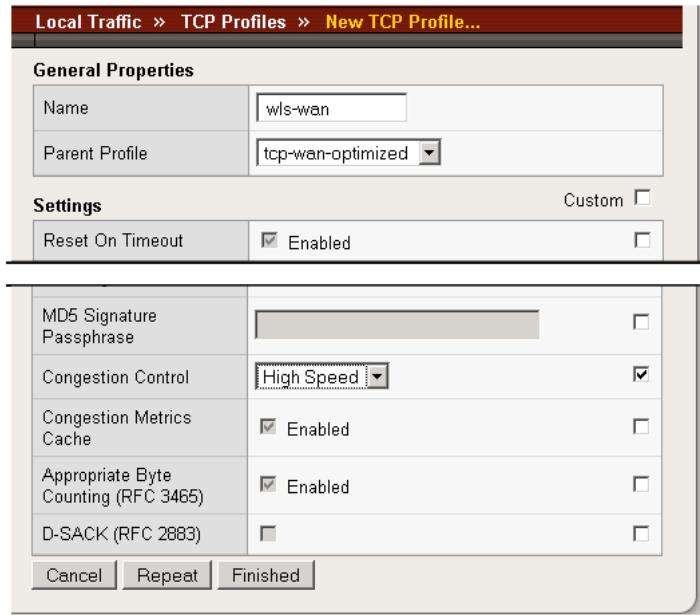


Figure 5 Creating the TCP profile (condensed to show relevant settings)

Creating the LAN optimized TCP profile

Next we configure the LAN optimized profile. Remember, if you are not using version 9.4 or do not want to use this optimized profile, you can choose the default TCP parent profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **wls-lan**.
5. From the **Parent Profile** list, select **tcp-lan-optimized** if you are using BIG-IP LTM version 9.4 or later; otherwise select **tcp**.
6. In the **Congestion Control** row, click the **Custom** box. From the list, select **High Speed**.
7. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.

Creating persistence profile

The next profile we create is a Persistence profile. We recommend using persistence for BEA WebLogic Servers, although the type of persistence depends on your configuration. In our example, use cookie persistence (HTTP cookie insert).

To create a new cookie persistence profile based on the default profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
The HTTP Profiles screen opens.
2. On the Menu bar, click **Persistence**.
The Persistence Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button.
The New Persistence Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **wls-cookie**.
5. From the **Persistence Type** list, select **Cookie**.
The configuration options for cookie persistence appear.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating a OneConnect profile

The next profile we create is a OneConnect profile. With OneConnect enabled, client requests can utilize existing, server-side connections, thus reducing the number of server-side connections that a server must open to service those requests. This can provide significant performance improvements for WebLogic implementations. For more information on OneConnect, see the BIG-IP LTM documentation.

In our example, we leave all the options at their default settings. You can configure these options as appropriate for your network.

To create a new OneConnect profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
The HTTP Profiles screen opens.
2. On the Menu bar, from the **Other** menu, click **OneConnect**.
The Persistence Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button.
The New HTTP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **wls-oneconnect**.
5. From the **Parent Profile** list, ensure that **oneconnect** is selected.

6. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the Stream profile (optional)

The next profile we create is a Stream profile. This profile is *optional*, and is created to correct instances where the web server inserts its own host name instead of the host name of the virtual servers, or incorrect paths into the content of the web pages.

The Stream profile performs a search and replace procedure for all occurrences of a string in a data stream efficiently and with minimal buffering. For more information on the Stream Profile, see *Solution 8115, Overview of the Stream Profile, on Ask F5*.

This procedure uses the host name of the virtual server you will create in *Creating the virtual server*, on page 14. If you do not yet know the host name, you can return to this procedure after creating the virtual server and modify the Target.

To create a new Stream profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
The HTTP Profiles screen opens.
2. On the Menu bar, click **Stream**.
The Stream Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button.
The New Stream Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **wls-stream**
5. Click the **Custom** box in the Target row. In the **Target** box, type use the following syntax:

`@<search>@<replace>@@<search>@<replace>@`

In our example, we type:

`@web0.weblogic10.tc.f5net.com:7041@web.weblogic10.tc.f5net.com@web1.weblogic10.tc.f5net.com:7041@web.weblogic10.tc.f5net.com@`

In this example, we are searching for the host name of the WebLogic Server, and replacing it with the host name of the virtual server you will create in *Creating the virtual server*, on page 14. The second search and replace pattern (following the @@) is for our second WebLogic Server.

6. Click the **Finished** button.

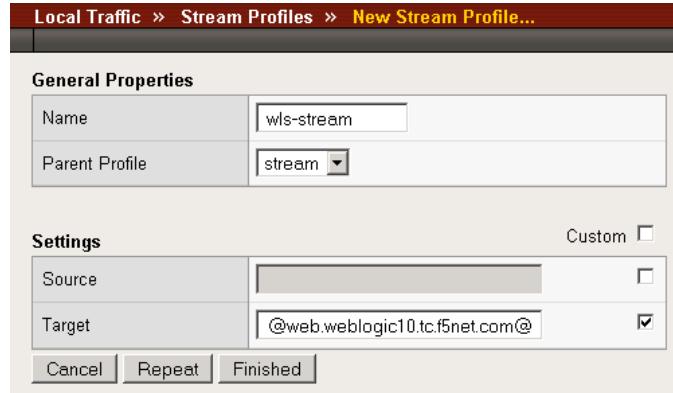


Figure 6 Creating the Stream profile

Creating a Client SSL profile

The next step in this configuration is to create a Client SSL profile. This profile contains the SSL certificate and Key information for offloading the SSL traffic. If you are not using the BIG-IP LTM system to offload SSL transactions, you do not need to create this profile.

To create a new Client SSL profile

1. On the Main tab, expand **Local Traffic**.
2. Click **Profiles**.
The HTTP Profiles screen opens.
3. On the Menu bar, from the SSL menu, select **Client**.
The Client SSL Profiles screen opens.
4. In the upper right portion of the screen, click the **Create** button.
The New Client SSL Profile screen opens.
5. In the **Name** box, type a name for this profile. In our example, we type **wls-clientssl**.
6. In the Configuration section, check the **Certificate** and **Key Custom** boxes.
7. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.
8. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
9. Click the **Finished** button.

Creating the virtual server

Next, we configure a virtual server that references the profiles and pool you created in the preceding procedures.

To create the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. In the upper right portion of the screen, click the **Create** button.
3. In the **Name** box, type a name for this virtual server. In our example, we type **weblogic-server**.
4. In the **Destination** section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **192.168.10.100**.
6. In the **Service Port** box, type **443**.



Figure 7 Creating the WebLogic virtual server

7. From the Configuration list, select **Advanced**. The Advanced configuration options appear.
8. Leave the **Type** list at the default setting: **Standard**.
9. From the **Protocol Profile (Client)** list select the profile you created in *Creating the WAN optimized TCP profile*. If you did not create a WAN optimized profile, select the LAN optimized profile as in the following Step. In our example, we select **wls-wan**.
10. From the **Protocol Profile (Server)** list, select the profile you created in *Creating the LAN optimized TCP profile*. In our example, we select **wls-lan**.
11. From the **OneConnect Profile** list, select the profile you created in *Creating a OneConnect profile*. In our example, we select **wls-oneconnect**.
12. From the **HTTP Profile** list, select the profile you created in *Creating an HTTP profile*. In our example, we select **wl-http-opt**.

13. If you are using the BIG-IP LTM for offloading SSL, from the **SSL Profile (Client)** list, select the profile you created in *Creating a Client SSL profile*. In our example, we type **wls-clientssl**.
14. Optional: If you created a Stream profile, from the **Stream Profile** list, select the profile you created in *Creating the Stream profile (optional)*. In our example, we type **wls-stream**.

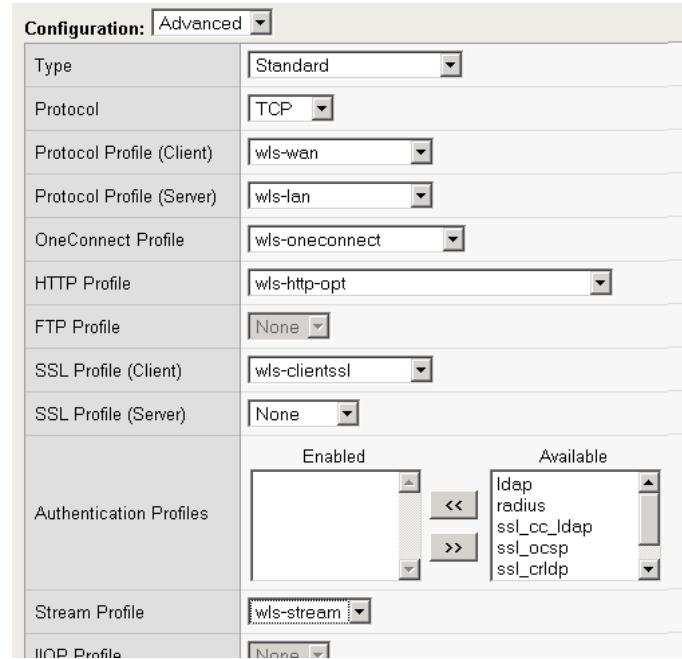


Figure 8 Selecting the BEA WebLogic profiles for the virtual server

15. In the Resources section, from the **Default Pool** list, select the pool you created in the *Creating the pool* section. In our example, we select **weblogic-server**.

16. From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating persistence profile* section. In our example, we select **wls-cookie**.

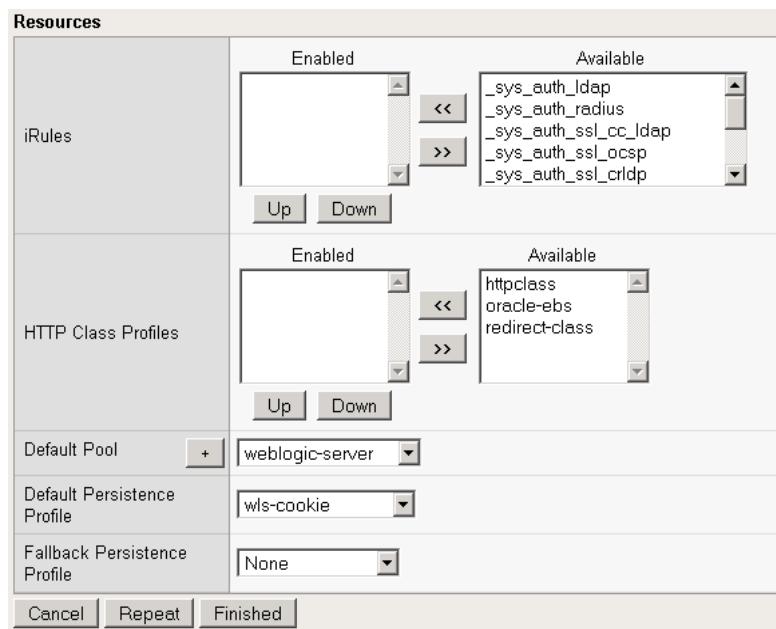


Figure 9 Adding the Pool and Persistence profile to the virtual server

17. Click the **Finished** button.

See the following section for WebLogic Portal configuration.

Configuring the BIG-IP LTM system for WebLogic Portal

In this section, we configure the BIG-IP LTM for BEA WebLogic Portal. Because the BIG-IP LTM configuration for WebLogic applications is similar to the Portal configuration, this section refers to procedures in the Portal configuration, with notes about any differences.

In many cases, you can use the same objects for both WebLogic Server and Portal (such as the health monitor and profiles), however, we strongly recommend you create new objects for each WebLogic component.

Creating a health monitor

The first task is to create a health monitor for the Portal devices. Follow the procedure *Creating the health monitor*, on page 3. Use a unique name for the monitor, and use the same a 1:3 +1 ratio between the interval and the timeout. All other settings are optional, configure as applicable for your configuration.

Creating the pool

The next task is to create a pool for the WebLogic Portal devices. Follow the procedure *Creating the pool*, on page 4. Type a unique name for the pool and configure the pool to use the health monitor you just created. Be sure to use the appropriate WebLogic Portal port (7041 by default) and IP address. All other settings are optional, configure as applicable for your configuration.

Using SSL certificates and keys

The next task is to ensure you have the appropriate SSL certificate and key on the BIG-IP LTM. Follow the procedure *Using SSL certificates and keys*, on page 6.

Creating the Profiles

For the WebLogic Portal configuration, you should create new versions of the profiles you created for WebLogic Server. With the exception of the Stream profile and the SSL profile, you could use the same profiles you created for the WebLogic Server configuration, but we strongly recommend you create new profiles. By creating new profiles, it makes it much easier to fine tune optimization and other settings for specific applications in the future.

For the HTTP profile, follow the procedure *Creating an HTTP profile*, on page 8, but in step 6, enter the appropriate Portal URIs.

For the optional Stream profile, follow the procedure *Creating the Stream profile (optional)*, on page 12, making sure in step 5 to use the appropriate host name and virtual server host name.

For the SSL profile, if you are importing a new certificate and key, follow *Importing keys and certificates*, on page 7, and use that information in the Client SSL profile.

Creating the Portal virtual server

The final step is to create a virtual server for the WebLogic Portal devices. Follow the procedure *Creating the virtual server*, on page 14. Give this virtual server a unique name, and use the appropriate address and port (443), and configure the virtual server to use all of the objects you created in the preceding procedures.

Modifying the WebLogic configuration

Once the BIG-IP LTM configuration is complete, you need to configure a Virtual Host on the WebLogic server that will use the LTM virtual server host name and IP address.

Configuring the WebLogic Server is outside the scope of this document, see http://edocs.bea.com/wls/docs100/ConsoleHelp/taskhelp/virtual_hosts/VirtualHosts.html for configuration information.

Make sure to use the fully qualified domain name (FQDN) of the BIG-IP LTM virtual server when creating the WebLogic Virtual Host.

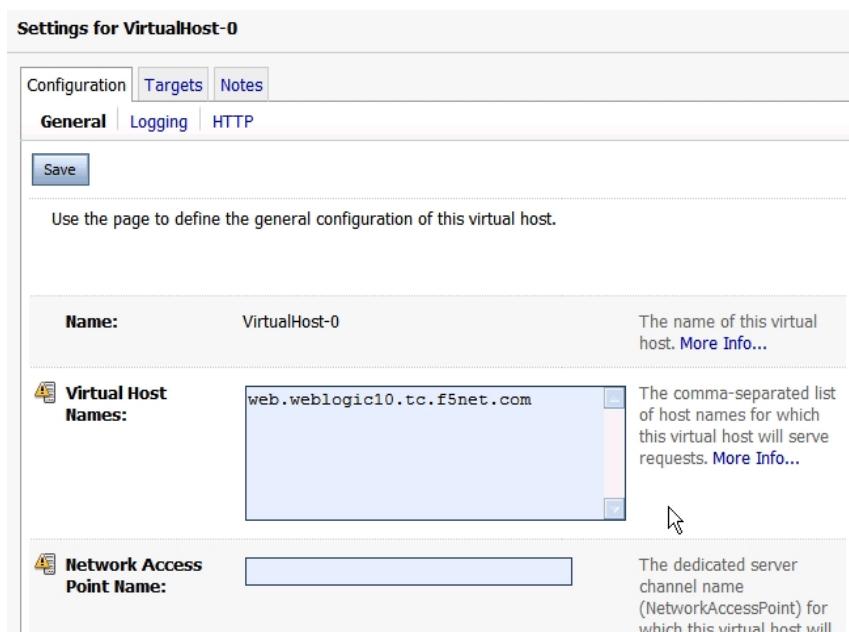


Figure 10 Configuring the WebLogic virtual host with the BIG-IP LTM virtual server host name

Synchronizing the BIG-IP LTM configuration if using a redundant system

If you are using a redundant BIG-IP configuration, the final step is to synchronize the configuration to the peer BIG-IP device.

To synchronize the configuration using the Configuration utility

1. On the Main tab, expand **System**.
2. Click **High Availability**.
The Redundancy screen opens.
3. On the Menu bar, click **ConfigSync**.
4. Click the **Self --> Peer** button.

Appendix A: Configuring the F5 WebAccelerator module with BEA WebLogic Server

In this section, we configure the WebAccelerator module for the WebLogic devices to increase performance for end users. The BIG-IP WebAccelerator is an advanced web application delivery solution that provides a series of intelligent technologies designed to overcome problems with browsers, web application platforms and WAN latency issues which impact user performance.

For more information on the F5 WebAccelerator, see
www.f5.com/products/big-ip/product-modules/webaccelerator.html.

Prerequisites and configuration notes

The following are prerequisites for this section:

- ◆ We assume that you have already configured the BIG-IP LTM system for directing traffic to the BEA deployment as described in this Deployment Guide.
- ◆ You must have purchased and licensed the WebAccelerator module on the BIG-IP LTM system, version 9.4 or later.
- ◆ This document is written with the assumption that you are familiar with the BIG-IP LTM system, WebAccelerator and BEA WebLogic Server. Consult the appropriate documentation for detailed information.

Configuration example

Using the configuration in this section, the BIG-IP LTM system with WebAccelerator module is optimally configured to accelerate traffic to BEA WebLogic servers. The BIG-IP LTM with WebAccelerator module both increases end user performance as well as offloads the servers from serving repetitive and duplicate content.

In this configuration, a remote client with WAN latency accesses a BEA WebLogic server via the WebAccelerator. The user's request is accelerated on repeat visits by the WebAccelerator instructing the browser to use the dynamic or static object that is stored in its local cache. Additionally, dynamic and static objects are cached at the WebAccelerator so that they can be served quickly without requiring the server to re-serve the same objects.

Configuring the WebAccelerator module

Configuring the WebAccelerator module requires creating an HTTP class profile and creating an Application. The WebAccelerator device has a large number of other features and options for fine tuning performance gains, see the *WebAccelerator Administrator Guide* for more information.

Connecting to the BIG-IP LTM device

Use the following procedure to access the BIG-IP LTM system's web-based Configuration utility using a web browser.

To connect to the BIG-IP LTM system using the Configuration utility

1. In a browser, type the following URL:
https://<administrative IP address of the BIG-IP device>
A Security Alert dialog box appears, click **Yes**.
The authorization dialog box appears.
2. Type your user name and password, and click **OK**.
The Welcome screen opens.

Creating an HTTP Class profile

The first procedure is to create an HTTP class profile. When incoming HTTP traffic matches the criteria you specify in the WebAccelerator class, the system diverts the traffic through this class. In the following example, we create a new HTTP class profile, based on the default profile.

To create a new HTTP class profile

1. On the Main tab, expand **WebAccelerator**, and then click **Classes**.
The HTTP Class Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button.
The New HTTP Class Profile screen opens.
3. In the **Name** box, type a name for this Class. In our example, we type **bea-class**.
4. From the Parent Profile list, make sure **httpclass** is selected.
5. In the Configuration section, from the **WebAccelerator** row, make sure **Enabled** is selected.
6. In the Hosts row, from the list select **Match Only**. The Host List options appear.
 - a) In the **Host** box, type the host name that your end users use to access the BEA WebLogic devices. In our example, we type **beaapplication.f5.com**(see Figure 11).

b) Leave the Entry Type at **Pattern String**.

c) Click the **Add** button.

d) Repeat these sub-steps for any other host names users might use to access the BEA deployment.

7. The rest of the settings are optional, configure them as applicable for your deployment.

8. Click the **Finished** button. The new HTTP class is added to the list.

Figure 11 Creating a new HTTP Class profile

Modifying the Virtual Server to use the Class profile

The next step is to modify the virtual server for your WebLogic deployment on the BIG-IP LTM system to use the HTTP Class profile you just created.

To modify the Virtual Server to use the Class profile

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.

2. From the **Virtual Server** list, click the name of the virtual server you created for the WebLogic servers. In our example, we click **bea-weblogic-vs**.
The General Properties screen for the Virtual Server opens.
3. On the Menu bar, click **Resources**.
The Resources screen for the Virtual Server opens.
4. In the HTTP Class Profiles section, click the **Manage** button.
5. From the **Available** list, select the name of the HTTP Class Profile you created in the preceding procedure, and click the **Add (<<)** button to move it to the **Enabled** box. In our example, we select **bea-class** (see Figure 12).
6. Click the **Finished** button. The HTTP Class Profile is now associated with the Virtual Server.

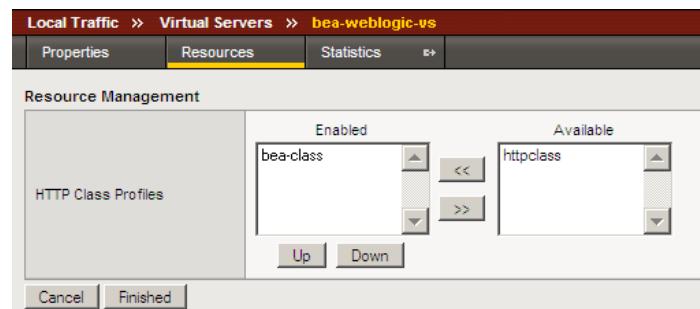


Figure 12 Adding the HTTP Class to the Virtual Server

◆ Important

If you are using the BIG-IP LTM version 9.4.2 or later, you must have created an HTTP profile on the BIG-IP LTM system that has RAM Cache enabled. In our example (Creating an HTTP profile, on page 8) we use a parent profile that includes RAM Cache. If you did not create an HTTP profile with RAM Cache enabled, you must create a new HTTP profile, based on a parent profile that uses RAM Cache (such as HTTP Acceleration), and modify the virtual server to use this new profile. This is only required for BIG-IP LTM version 9.4.2 and later.

*To create the HTTP profile, use **Creating an HTTP profile**, on page 8, selecting the HTTP Acceleration parent profile. You must leave RAM Cache enabled; all other settings are optional. To modify the virtual server, follow Steps 1 and 2 from the preceding procedure to access the virtual server, and then from the HTTP Profile list, select the name of the new profile you just created and click Update.*

Creating an Application

The next procedure is to create a WebAccelerator Application. The Application provides key information to the WebAccelerator so that it can handle requests to your application appropriately.

To create a new Application

1. On the Main tab, expand **WebAccelerator**, and then click **Applications**.
The Application screen of the WebAccelerator UI opens in a new window.
2. Click the **New Application** button.
3. In the Application Name box, type a name for your application. In our example, we type **BEA WebLogic**.
4. In the **Description** box, you can optionally type a description for this application.
5. From the **Local Policies** list, select **BEA WebLogic**. This is a pre-defined policy created specifically for BEA WebLogic devices (see Figure 13).
6. In the **Requested Host** box, type the host name that your end users use to access the BEA WebLogic deployment. This should be the same host name you used in Step 6a in the preceding procedure. In our example, we type **beaapplication.f5.com/**.
If you have additional host names, click the **Add Host** button and enter the host name(s).
7. Click the **Save** button.

General Options	
Application Name:	BEA WebLogic
Description: (optional)	This is a WebAccelerator application for WebLogic

Policies	
Central Policy:	BEA Weblogic
Remote Policy:	- Select One -

Hosts	
Requested Host	Action
iis-application.f5.com	Options Delete

Add Host **Save** **Cancel**

Figure 13 Configuring an Application on the WebAccelerator

The rest of the configuration options on the WebAccelerator are optional, configure these as applicable for your network. With this base configuration, your end users will notice a marked improvement in performance after their first visit.